

Unconfirmed peers and spinelessness

Ben Sherman

Department of Philosophy, Boston University, Boston, MA, USA

ABSTRACT

The Equal Weight View holds that, when we discover we disagree with an epistemic peer, we should give our peer's judgment as much weight as our own. But how should we respond when we cannot tell whether those who disagree with us are our epistemic peers? I argue for a position I will call the Earn-a-Spine View. According to this view, parties to a disagreement can remain confident, at least in some situations, by finding justifiable reasons to think their opponents are less credible than themselves, even if those reasons are justifiable only because they lack information about their opponents.

ARTICLE HISTORY Received 28 July 2014; Accepted 30 August 2015

KEYWORDS Epistemology; disagreement; epistemic peer; Equal Weight View; spinelessness

One prominent view about disagreement, sometimes called the *Equal Weight View* (EWW), argues that, when we discover we disagree with one or more epistemic peers, we should give each peer's judgment as much weight as our own (and likewise for any peers that agree with us). There are powerful and elegant arguments for this view, but it seems to have alarming implications for our controversial beliefs on matters like religion, ethics, politics, metaphysics, and so on. After all, is it not the case that many of those who disagree with us on these topics are our epistemic peers? If so, wouldn't the EWW require us to suspend judgment, or at least become fairly wishy-washy, about many of our most cherished beliefs? Adam Elga refers to this seeming implication as *spinelessness* (2007, 484).

But spinelessness does not follow directly from the EWW, since it is possible that few – or none – of those who disagree with us are, in fact, our epistemic peers. Moreover, even if some (or most, or all) of those who disagree with us are our peers in fact, we often have too little information to tell whether or not they are our peers. The EWW only tells us what to do in response to *recognized*

CONTACT Ben Sherman  bens@bu.edu

¹ Shortly prior to this article's acceptance, the author accepted a position at Brandeis University.

© 2015 Canadian Journal of Philosophy

peer disagreement. As Christensen (2011, 15–16) points out, the claim that peer disagreement gives us reason to revise our beliefs does not entail that we must *presume* that those who disagree with us are our peers when there is room for doubt. And there is certainly plenty of room for doubt about who (if anyone) is an epistemic peer on matters like religion, ethics, politics, etc.

So the EWV, insofar as it only addresses recognized peer disagreement, does not necessarily commit us to spinelessness. But it seems doubtful that anyone who defends the EWV intends it to be so narrow a principle; rather, it is suggestive of a more open-minded and intellectually humble approach to disagreement. As King (2012, 267–269) points out, the fact that someone is not *known* to be an epistemic peer does not entail that we are *not* obliged to revise our beliefs when they disagree with us. So, for those of us who find the EWV fairly convincing, and want to figure out how to react to disagreements about religion, philosophy, etc., the most pressing question seems to be: How should we respond to disagreement with *unconfirmed peers* – that is, those whom we can neither identify as peers nor rule out as peers?

The EWV itself is controversial, and I will not attempt to defend the view against its critics here. My project here is rather to consider how best to extend the EWV to disagreement with unconfirmed peers. Sections I and II offer some clarification of the concept of epistemic peerhood, and respond to the view that we are unlikely to have any epistemic peers on topics like philosophy and religion. Sections III and IV review and criticize several approaches to disagreement with unconfirmed peers available in the literature. In Sections V–VIII I propose and defend my own position on disagreement with unconfirmed peers, which I call the Earn-a-Spine View.

I. Do we really face disagreement with epistemic peers?

King (2012, §1) argues that, when it comes to matters of any complexity, we are unlikely to have any epistemic peers at all – after all, it is rare that two people have the very same evidence and reasoning capacities. If King is right, we might be able to neatly sidestep the whole problem of peer disagreement. But matters are not quite so simple; King's argument turns on a conception of the term 'epistemic peer' subtly at odds with some of the central thinkers in the literature.

According to King's interpretation of the literature on peer disagreement, it is generally accepted that, for two people, S and T, to be epistemic peers with regard to a proposition, P, they must meet the following conditions (along with some others that are less troublesome):

The same-evidence condition: S and T have the same P-relevant evidence, E. *The dispositional condition:* S and T are equally disposed to respond to E in an epistemically appropriate way. (252)

But King underestimates how diverse notions of epistemic peerhood are.¹

The trouble is that various epistemologists' definitions of peerhood have differed in at least three important ways. Every notion of peerhood in some way tries to describe peers as comparably credible, and the differences between various definitions involve differences in the credibility-conferring features stipulated:

- (1) Different definitions of peerhood sometimes have different items on their lists of credibility-conferring features. (These differences can be interesting, but I will largely ignore this kind of difference between definitions.)
- (2) Some thinkers define epistemic peers as those with *the same* credibility-conferring features (e.g. Elgin 2010, 53), while others define epistemic peers as those with *equally good* credibility-conferring features (e.g. Gutting 1982, 83; Kelly 2005, 174–5; Elga 2007, 484; Christensen 2009, 756–757; Conee 2009, 313; Lackey 2010, 203).²
- (3) Some thinkers define epistemic peers as being equal (or the same) in *each* of various credibility-conferring features (e.g. Gutting 1982, 83; Elgin 2010, 53; Lackey 2010, 203), while others define peers as being equal in terms of their *overall set* of credibility-conferring features (e.g. Elga 2007, 484; Conee 2009, 313).³ Some thinkers' definitions are ambiguous in this respect (e.g. Kelly 2005, 175; Feldman 2007, 201; Christensen 2009, 756–757).

King argues, correctly I think, that if peerhood involves having *the same* P-relevant evidence, we will almost never encounter epistemic peers, at least not on matters of any complexity (2012, §1.2; cf. also Lackey [2010, 311]). Two experts in the same field will have read some different studies and arguments. People's experiences will make it reasonable to give more or less credence to various authorities. If first-person intuitions or religious experiences qualify as evidence, two people certainly cannot share this evidence. Those who define epistemic peers as having *the same* epistemic dispositions present a situation we are even less likely to encounter.

It is much less surprising to encounter someone with evidence just as good as one's own, or a set of epistemic dispositions that, taken as a whole, are just as good. They might not be *exactly* equal in either of these respects, any more than we are likely to encounter people who are *exactly* the same height – but we will certainly encounter people who are so close to equal that the difference is trivial, and sometimes undetectable. Two experts may have read some different studies, or had slightly different intuitions, but that should not lead us to conclude one expert's evidence that P is any greater than the other's; they may be different, but equally good. Likewise, they may have different, but equally good training, or one might have more experience while the other is quicker-thinking.

Moreover, it is not hard to imagine two people being in equally good positions to determine whether P is true without being equal in evidence or dispositions;

after all, one might have more salient evidence, and the other might have better dispositions for responding to the evidence.

King (2012, 268–269) makes a similar point, suggesting that ‘one’s *total epistemic position*’ is more worthy of investigation than peerhood. My dispute with King, then, is largely semantic. While I agree with King that, as *many* thinkers define epistemic peerhood, it is something of a red herring, Elga is a clear exception; by his definition, you regard someone as your peer, with respect to a certain sort of claims, if you regard her ‘as being as good as you at evaluating such claims’ (Elga 2007, 484). Since Elga also coined the widely-used term *EWV*, and identified the problem of spinelessness, I propose we take his definition of *epistemic peer* to be the one most appropriate to discussions about whether the EWV demands spinelessness.⁴

Pedigree aside, Elga’s definition seems most appropriate for a couple of reasons. First, if I recognize that someone who disagrees with me is *just as likely to be right* as myself, there is no obvious justification for giving their opinion less weight than that of someone who has *the same* evidence and dispositions I do – so if we focus only on those with the same evidence and dispositions, we are unduly narrowing the scope of the discussion. Second, while there is probably no one who has even roughly the same evidence I have for my ethical, political, and philosophical views, I suspect that there are a great many people who have a (more or less) equally good overall set of credibility-conferring features, and so are (more or less) equally likely to be right. Since it is very likely I have Elga-type peers, it makes sense to worry about whether those who disagree with my philosophical, political, and religious views are my Elga-type peers.

Why all the focus on epistemic peers? They seem to be the focus of the debate, in part, because they prompt differing intuitions from different camps. Opponents of the EWV sometimes argue, for instance, that it can sometimes be reasonable not to give *any* weight to an epistemic peer’s judgment – that is, to remain just as confident of an opinion as you were prior to discovering the disagreement (cf. Kelly 2005). The EWV is distinguished, of course, by its position that one should give equal weight to the view of every epistemic peer – that is, treat every epistemic peer as equally likely to be right, and so update one’s belief when one discovers that epistemic peers disagree. Those (like Feldman) who take an all-or-nothing view of belief generally hold that we should suspend judgment if a significant portion of epistemic peers are divided over an issue, while those (like Elga and Christensen) who think of beliefs in terms of probabilities argue that we should average our credences with those of our epistemic peers (at least if no other opinions are relevant to the proposition in question.)

It should be noted that no version of the EWV says disagreement with an epistemic peer is the *only* kind of disagreement that should lead us to revise our beliefs. All versions of the EWV (and some views opposed to it) agree that an epistemic *superior’s* judgment should have even more weight than one’s own. Many versions of the view also hold that we should sometimes revise our views

somewhat in the face of disagreement with epistemic *inferiors* – for instance by giving their view some weight, but less than one’s own – as long as the inferior in question has at least some credibility. But peers remain interesting at least partly because, in deep disagreements, we usually cannot identify those who disagree with us as clearly superiors or inferiors.

II. Unconfirmed peers

Arguments in favor of the EWW often appeal to thought experiments in which two people discover a disagreement about some perceptual judgment (cf. Feldman 2006, 223) or fairly simple arithmetical calculation (cf. Christensen 2007, 193). In these cases, it is highly plausible that the parties to the disagreement should become much less confident; there are only two parties whose judgments are relevant, and we can easily suppose that the two parties have ample reason to suppose their evidence and capacities are equally suited to the matter under dispute.

Philosophical, political, and religious disagreements are clearly another matter. In these cases, we do not have ample reason to suppose those who disagree with us are our peers. Yet, by the same token, we don’t have sufficient reason to decide they are *not* our peers. In such debates, we frequently face disagreement with *unconfirmed peers*. In a widely debated disagreement, we can reasonably expect that at least a few of our epistemic peers or superiors disagree with us on any given disputed proposition, but we will often lack information about whether the majority of the most qualified opinions favor one side or the other. As long as *most* of those who disagree with us are unconfirmed peers, it is unclear how the EWW would have us react to the situation.

III. How to respond to unconfirmed peers – two positions

In a somewhat infamous passage, Elga (2007, 492–497) argues that the EWW will not lead to spinelessness. He presents a thought experiment in which Ann and Beth, two friends at opposite ends of the political spectrum, disagree about whether abortion is morally permissible, and he argues that neither party needs to significantly reduce her confidence, because each thinks the other is mistaken about a whole range of associated questions, and hence is less than an epistemic peer (493). But then, of course, it makes sense to stop looking at the abortion debate in isolation, and consider the broader debate: can Ann and Beth both rationally think the other less than an epistemic peer if neither can give a non-question-begging defense of her whole network of controversial views? Elga seems to think they can, and offers a defense of Ann’s doing so (which could equally well apply to Beth):

Consider the cluster of issues linked to abortion ... setting aside her reasoning about the issues in the cluster, and setting aside Beth’s opinions about those issues,

Ann does not think Beth would be just as likely as her to get things right. That is because there is no fact of the matter about Ann's opinion of Beth, once so many of Ann's considerations have been set aside. (495–496)

But even if we did have so little common ground with others that even our epistemic standards are called into question, the fact that we cannot offer a non-question-begging defense of our own network of assumptions seems to show that we have no reason we can offer for thinking we are right and the others are wrong. Elga moves too quickly from a situation where we cannot tell who is better justified to the conclusion that we can remain confident that the opposing view is not justified.

Foley (2001) also makes the case that it is rational to remain confident in such situations. Following Gibbard (1990, 180–182), he argues that each individual must take a 'leap of faith' in trusting her own judgment, and this leap of faith is necessary to escape bleak skepticism (Foley 2001, ch. 1.6). Because self-trust is of such fundamental importance, Foley favors what I will call *The Presumption in Favor of Self-Trust*: we should trust our own judgment *unless we have reason to think others' judgment is as good or better*.⁵

Foley's argument for the presumption in favor of self-trust is motivated by the argument that without self-trust we cannot avoid bleak skepticism. But this argument by itself does not show that we must prioritize our standards over others (cf. Feldman 2006, 224). Bleak skepticism can be avoided so long as we trust people's cognitive systems in general, including, but not especially, our own. Catherine Elgin argues that we are only *justified* in thinking our powers of judgment are reliable if others mostly corroborate our judgments upon reflection (1996, 116). As we rely on others to corroborate and inform our own beliefs, we are obliged 'to treat our compatriots' commitments as we do our own' (118). A plausible interpretation of Elgin's view (and Feldman's [2006, 2007]) is that they accept what I will call *The Presumption of Peerhood*: when we know of disagreement, we should presume others are our epistemic peers, until we find mutually recognizable evidence of epistemic superiority on one side or the other.

But the Presumption in Favor of Self-Trust and the Presumption of Peerhood have a problem in common: they both recommend that, when in doubt, we make a certain kind of presumption, pending new information. Either way, we can expect the presumption to be wrong a non-trivial portion of the time. The Presumption in Favor of Self-Trust will leave us frequently underestimating others, while the Presumption of Peerhood will leave us sometimes overestimating others and sometimes underestimating others. If we could do no better than leaping to a conclusion and being prepared to revise it, this would not be so bad; it could be a matter of taste or circumstance which presumption we followed. But can we do better?

IV. The mushy approach

While Feldman takes an ‘all-or-nothing’ view of belief (such that our only options for a given proposition are to believe it, deny it, or suspend judgment), most others in the debate think belief can be a matter of degree, and can be represented as probabilities, or credences. When we face disagreement with an unconfirmed peer, how should our credences change? It seems the presumption in favor of self-trust would have us make the minimum adjustment to our credences consistent with our evidence about the credibility of those who disagree with us. The presumption of peerhood seems to suggest that, pending evidence of inequality, we average our credences with those of unconfirmed peers. But, again, each of these approaches helps itself to information we don’t have. Since we are uncertain about the credibility of an unconfirmed peer, it seems precisely the question at issue is how much we should adjust our credences.

This sort of higher order uncertainty is exactly the sort of problem that has motivated some thinkers to propose that sometimes it is more rational to have ‘mushy’ or ‘indeterminate’ credences. Instead of averaging my credences with an unconfirmed peer, my credences might cover a range, with the lower bound of this range representing the credence I would have if my unconfirmed peer has the highest credibility that seems plausible to me, and the upper bound being the credence I would have if she had the lowest credibility that seems plausible to me.

The Mushy Response might be the theoretically ideal response to disagreement with unconfirmed peers, but I will seek another approach, for a couple of reasons. First, mushy credences are a controversial topic at the moment, and it might turn out that they are not theoretically cogent.⁶ Second, I would prefer to develop a response to unconfirmed peers that can be applied even by those who take an all-or-nothing view of belief. So, I will propose a response that attempts to be neutral between different conceptions of belief.

V. The Earn-a-Spine View

My view, which I will call the *Earn-a-Spine View* is this: when you encounter disagreement with one or more unconfirmed peers, your response to the situation depends on your answers to two questions:

Question 1: Do you think there is an epistemic asymmetry between you? (Do you think it is more likely that you are right? Do you think it is more likely that they are right?) If not, then you should regard the unconfirmed peer as an epistemic peer.

If you do think there is an asymmetry, then ...

Question 2: What do you think is the source of the epistemic disparity? (What do you think makes it more likely that you will be right? Or that they will be right?)

If you have no answer to *Question 2*, or you are not justified in accepting your own answer to *Question 2*, you should regard the unconfirmed peer as an epistemic peer.

If you have an answer to *Question 2*, and you are justified in thinking it is right, then, to the extent that answer gives you reason to think there is an epistemic asymmetry, you can give unequal weight to the unconfirmed peer's view.

VI. The argument for the Earn-a-Spine View

When we encounter disagreement with unconfirmed peers, we are uncertain about how much weight to give their judgment because the evidence we have for our own credibility is not readily comparable to the evidence we have (if any) of their credibility. In most situations, we will have much more evidence about our own credibility than we will have about others' (though this could be reversed in rare cases.) I have argued that both the Presumption in Favor of Self-Trust and the Presumption of Peerhood leap too readily to unsupported hypotheses. But are there situations in which it is reasonable to remain confident in our own view, despite encountering disagreement with an unconfirmed peer?

Christensen makes a promising observation:

Typically, when I am highly confident in my initial opinion, I have good reason to think that the opinion is based on highly reliable reasoning. But this itself gives me some reason to think that an equally informed person who disagrees with me did not use the same sort of reasoning I did, since it is unlikely that two people, using a highly reliable method of reasoning on the same evidence, would reach different opinions. So in many cases where I know relatively little about the person with whom I disagree, my having a great deal of confidence in my initial opinion *should* correlate with my giving less credence to the opinion of the other person. (2007, 203)

I think Christensen's point about methods of reasoning can be generalized to other credibility-conferring features, including evidential bases for belief. If some passerby on the street tells me that Socrates wrote the *Iliad*, it is possible the person is a brilliant classicist who has recently unearthed shocking new information, but it seems more likely to me that the person is just not very familiar with ancient Greek figures, and is confused. This does not mean I am presuming I am more credible by default; rather, my philosophical training has made me much more familiar with Socrates than the average person on the street, and it seems reasonable to doubt that this person is an exception, pending evidence to the contrary.

My putative explanation of the disagreement does important work in this scenario. I cannot just presume I have better credentials than the person who disagrees with me; I must have some further opinion about *which* credentials were lacking. This extra opinion is important, as it is now open for challenge and revision, at least in principle. If I later conclude that the opinion is implausible (if, for instance, I later find out that the person on the street *was* a leading

classicist) I must also revisit my conclusions about the disagreement. Having an opinion about *why* we think others are not our peers makes us more intellectually accountable. In this way, we might be able to remain confident about (at least some of) our cherished beliefs, even while those who disagree with us remain confident of their own cherished beliefs, but all have to earn the privilege by coming up with an explanation of the disagreement that justifies their confidence. By identifying a basis for remaining confident, we 'earn a spine.'

Likewise, we might suspect our unconfirmed peers are *more* credible than ourselves. In those cases, similar standards apply: we must identify some reason for thinking there is such an asymmetry, and that reason must meet our intellectual standards; and then we should adjust our beliefs as is appropriate for the asymmetry we have decided we face. I expect that this situation is less common than the one in which we think ourselves more credible than unconfirmed peers, so I will mostly focus on the latter type.

What if you have no explanation for the disagreement in mind, and have too little evidence to decide whether we are epistemic peers? Should you give my view equal weight? Maybe not, if you have the option of making your credences mushier. But if that is not an option, I think you should give my view equal weight to your own.

Why should you assume I am your equal on this matter, when you have too little evidence to decide whether this is true?

The quick and dirty answer is that doing so is a way to hold yourself accountable for your beliefs. Like most people, you are biased in your own favor; if, in spite of your biases, you are not able to find *any reason at all* to think I am more likely to be wrong, that is bad news for you. Since you are highly attuned to evidence for your own position and your own credibility, this is a situation where your lack of evidence of an epistemic advantage is *prima facie* evidence of a lack of epistemic advantages.

The somewhat more technical answer is this: if you can suppose yourself to be my epistemic superior with regard to any proposition, *P*, on which we are unconfirmed peers, and about which you cannot think of any particular reason to believe you are my epistemic superior that enables you to engage in bootstrapping.

Suppose that, on any matter where we are unconfirmed peers, and where you cannot find any reason at all to think you are more likely to be right, you can suppose there is a 51% chance that you are right. Given the nature of the situation, you would have had no basis, before discovering we disagree about *P*, to think you would be at all more likely than me to be right about whether *P*. But, now that we have learned that we disagree, you can decide that you are slightly more likely than me to be right. But the disagreement itself is your only reason for thinking so, and the mere fact of disagreement reveals nothing about *which* of us is more likely to be right.

The problem becomes even more extensive if we discover a range of such disagreements. Then, every time we discover another disagreement, you gain some evidence you have a better track record than me. But it cannot be reasonable to grow increasingly confident you have a better track record than me just on the basis of finding that we disagree on a wide range of matters; you need to have a reason for thinking your judgments were more reliable or well founded. That is, you have boosted your grounds for thinking yourself more credible just by thinking yourself more credible – you have lifted yourself up ‘by your bootstraps.’⁷ A good epistemological position should avoid allowing this sort of bootstrapping, and here we can avoid it by considering each other epistemic peers until we can come up with some reason to think one party is more credible than the other.

But how often are you really in a situation where you can’t find *any reason at all* to think I’m more likely to be wrong? I suspect that, when we actually face such situations, we tend to find it natural and reasonable to doubt our views. But I might be wrong, and even if I’m right, I’m sure there are exceptions. Still, it seems that we usually think we *do* have some kind of epistemic advantage;⁸ in that case we think there is some reason to believe those who disagree with us are not generally our peers on the matters under dispute. If we are satisfied with the mere impression that we must enjoy *some* advantage or other, we are giving ourselves a license to bootstrap and failing to hold ourselves accountable for our beliefs.⁹ But that does not mean we must conclude we do *not* enjoy an epistemic advantage when we feel confident about our disputed beliefs. We avoid the problems of bootstrapping and unaccountability if we figure out just what epistemic advantage we think justifies our confidence, and examine whether we actually have good reason to believe in this epistemic advantage.

If you identify the epistemic advantages you think you enjoy, have you earned a spine? That depends. If you recognize that it is irrational or irresponsible to think you enjoy these advantages, they don’t justify confidence in the face of controversy, and you have not earned a spine. The EWV and Earn-a-Spine View are neutral between various standards of rationality and responsibility, so I will leave it an open question which standards are involved; and, of course, sometimes we disagree about standards of rationality and responsibility, in which case we would want these standards to be part of the debate. But identifying your putative epistemic advantage is an important necessary condition for earning a spine: when you identify the epistemic advantages you think you enjoy, your ideas are now subject to appraisal; they make you accountable; and (however flimsy they might be) they offer better support than your bootstraps. Disagreement alone cannot show that you are more likely to be right than your opponent, so a minimum condition for earning a spine is identifying what could make you more likely to be right.¹⁰

VII. How difficult is it to earn a spine?

Since I am not able to specify here a full theory of epistemic rationality and responsibility, and since these standards could be part of what is disputed in a given disagreement, I cannot give a full account of what is required in earning a spine, except to say that your explanation for your epistemic superiority must meet your own standards of rationality and responsibility.¹¹ There are a few conditions we can name right off the bat however, whatever other standards of rationality and responsibility should be applied:

- (1) You cannot appeal to an epistemic advantage that you know your opponent thinks you lack, unless you can base this appeal on evidence you reasonably think your opponent does not reject. (So you usually cannot get away with 'My opponent is too stupid to decide this sort of question, so I have an advantage!' unless you can back it up with something like '... and my opponent would agree with me if he knew as much as I do about the indicators of stupidity, and the ways stupidity makes people unable to decide this sort of question.')
- (2) You cannot appeal to an epistemic advantage you know your opponent denies is an epistemic advantage, unless you can base this appeal on evidence or standards you reasonably think your opponent does not reject. (So you cannot get away with 'I have an epistemic advantage over my opponent who thinks meditation is a waste of time, because I meditate, and that makes me better at figuring these sorts of things out' unless you can back it up with something like '... and my opponent would agree that this is an advantage if she thought about it, because surely she accepts that someone who is familiar with a practice is better at evaluating it.')
- (3) You cannot appeal to a suspected disadvantage on your opponent's part if you know your opponent has just as much reason to suspect you of the same disadvantage. (So you cannot get away with 'Maybe my opponent has a hard time accepting new ideas' without some comparative evidence. Although you might, in fact, be better at accepting new ideas than your opponent, the fact that this is a shortcoming we cannot identify through introspection means we cannot claim any sort of advantage merely on the basis of suspecting another has the shortcoming.)
- (4) You cannot claim new evidence or arguments as an advantage if you have reason to think your opponent is likely to have just as much new evidence or argumentation against your view. (So you cannot get away with 'Although our last debate ended in a draw, this time I surely have an advantage, because I have done further research!' unless you reasonably think your opponent has not been doing further research.)

These conditions are not trivial; it is easy to become overconfident by thinking about errors others might be making without thinking about the errors we might be making. Still, it should be clear that they permit us to think we enjoy an epistemic advantage despite having asymmetrical evidence about ourselves and our opponents. (1) and (2) forbid us to beg the question against our opponents, but leave open the possibility of thinking we enjoy an epistemic advantage in the disagreement because we underestimate the breadth of the disagreement. (3) and (4) rule out thinking we have an advantage merely on the basis of suspecting a shortcoming on our opponents' part, or knowing of new support for our own position, but they leave open the possibility of drawing conclusions about an advantage when we have far more information about one side than the other (that is, we have less evidence about our own blind spots and more evidence about our own defenses, than we do about our opponents' blind spots and defenses, respectively.) I suspect these evidential asymmetries are at work in many ongoing debates about cherished beliefs, especially in the academic world.

Of course, asymmetry of evidence alone does not justify us in thinking we enjoy an epistemic advantage; asymmetry of evidence is not necessarily evidence of asymmetry. When are we justified in concluding we have an advantage on the basis of asymmetrical evidence about ourselves and others? That will depend on our more general views on rationality and responsibility. But then, is it *ever* reasonable to be confident we enjoy an epistemic advantage when we have asymmetrical evidence about those who disagree with us?¹²

Here is one reason to think it is: as Gibbard (1990, 180), Elgin (1996, 116–118), and Foley (2001, ch. 4.1–4) argue, we cannot escape bleak skepticism without having some kind of trust in our shared human cognitive capacities. Trusting those capacities means having a sort of *prima facie* trust in whatever some human being – including yourself – thinks is true. If it seems to you that *P*, you are at least a little justified in thinking that *P*, unless you know that someone else thinks that $\sim P$, or you have some other sort of evidence that you cannot trust yourself to be right about judgments like *P*. So, if it seems to you that you have a certain epistemic advantage over me, and you have not heard anyone dispute that impression, and it is not a sort of impression you have learned is untrustworthy, then you are at least somewhat justified in thinking you are right. Some might say you should assume I would dispute your claim to this epistemic advantage, but this assumption would do me a disservice; if you have, in fact, noticed a real epistemic advantage you have over me, I would like to believe I would admit it when it was pointed out to me. (Of course, I might well fail to understand your point, or might unreasonably reject your claim; but, with luck, my incomprehension or unreasonableness would give you still more evidence of your epistemic advantage.) You must be prepared to reconsider your judgments if they are disputed, or otherwise fit poorly with new evidence you acquire. But

as long as your judgments go unchallenged, you are at least a little justified in thinking they are right.¹³

Other questions about how much to trust a new idea will depend on more general theories of rationality and responsibility. They will include questions about how to make use of information about our capacities for error, how our confidence should reflect our track record (or lack thereof), how we should account for unknown unknowns, and so on. But if the Gibbard–Elgin–Foley argument is right, at least a little *prima facie* confidence is warranted for anything you think is true.

The natural worry about the Earn-a-Spine View is that we will wind up being overconfident; even when those who disagree with me are, in fact, my epistemic peers or superiors, I will find reasons to think I have an epistemic advantage. But this is not that bad a thing for two reasons.

First, the Earn-a-Spine View does not give us license to reach conclusions we can recognize as irrational or irresponsible. So, the fact that I could falsely think myself epistemically superior to others is just an instance of the familiar fact that sometimes we can be justified in believing something that is false. This would only be a special problem for the Earn-a-Spine View if I could insulate my beliefs from challenge after having earned a spine. But, on the contrary, the Earn-a-Spine View holds us accountable for naming the epistemic advantages we think we enjoy, and doubting them as soon as they are challenged and we are unable to give a non-question-begging defense.

Second, the Earn-a-Spine View should make disagreement more productive in those situations where one party (or more) is overconfident. Suppose that you and I disagree about whether *P*; you think that *P*, and I think that $\sim P$. We find that neither of us has a non-question-begging argument for our views; you reject many of my premises, and I reject many of yours. Since neither of us can find a reason to think we have an advantage at the moment, our conversation ends with each of us taking the other to be an epistemic peer. But, after our conversation ends, each of us thinks about why the other seems mistaken. You think of several epistemic advantages you think you enjoy, that make you better qualified to decide whether *P* and the premises for and against it are true. I think of several epistemic advantages I think I enjoy. As it turns out, you are right about all your epistemic advantages, and I am wrong about all mine; you really were more qualified all along, though neither of us had figured that out at the time of our discussion. One benefit of the Earn-a-Spine View, then, is that it encouraged you to figure out what epistemic advantages you enjoyed.

But what about me? I now accept a network of mistaken views, and that network has grown since I came to falsely think I have various epistemic advantages over you. And as a result of having more false views, I am *more* confident. This may seem like a bad thing, but it's not. By doing the work of spelling out reasons, we both created networks of beliefs implicated in the disagreement. I do not just have more false beliefs; I have a network of false beliefs. The

work I have done in earning a spine pays off by showing how large a network needs to be reconsidered if I discover part of it is flawed. Given the assumption that creatures like us have some tendency to accept truth over falsity, a problem is more likely to be found (all else equal) in a broad network of false beliefs than in isolated false beliefs; if I realize some part of the network is mistaken, other parts of the network will be implicated and flagged for reconsideration.

Of course, there is also the possibility that I will never recognize a problem in my network of false views. This is sad, but it is by no means unique to the Earn-a-Spine View. The Presumption in Favor of Self-Trust makes this more likely, and the Presumption of Equality makes it more likely that a network of views favored by the majority will go uncorrected; the same goes all the more for views that deny we need to become less confident in the face of disagreement at all. The Earn-a-Spine View at least makes me work for it, and denies us the luxury we often enjoy, of trusting a feeling that we are better qualified than our opponents without holding ourselves accountable for spelling out our qualifications.

VIII. Is it too easy to earn a spine?

My proposal, in brief, is that we can, at least sometimes, take ourselves to be more credible than unconfirmed peers, on the basis of asymmetrical evidence: we know we have credibility-conferring features, and think our unconfirmed peers lack them, for instance. But, if we all want to remain confident in our cherished beliefs, does this standard end up encouraging bad epistemic behavior?

In some ways, it does encourage bad behavior. But I think the bad behavior it encourages is no worse than that encouraged by any other internalist position.

Suppose I know that many of my unconfirmed peers disagree with my political views. I want to earn a spine, so I consider epistemic advantages I might have over all (or at least most) of those who disagree with me. What is to stop me from thinking that all those who disagree with me have just failed to pay close attention to the issues?

If that thought really is plausible, given the evidence available to me, perhaps I can earn a spine that way. But, as is, that idea does not accord well with my evidence; it is clear to me that many people who disagree with me pay close attention to the issues – and some pay much closer attention than I do. So that thought would not meet my standards of rational or responsible acceptance.

But what if I could alter my epistemic situation, to make it rational and responsible to adopt that view? Perhaps I could somehow cause myself an epistemic defect, such that I forgot, or failed to notice, all the evidence showing that people who disagree with me pay close attention to the issues. The Earn-a-Spine View has the distasteful implication that I could achieve my goal of rationally believing that I was superior to my unconfirmed peers by diminishing my own epistemic state. But note that the same problem appears for any internalist theory of epistemic justification. If I want my belief, *P*, to be rational, but it does

not accord with my evidence, evidentialists must admit that I could achieve my goal by causing myself to lose the evidence that conflicts with my target belief.¹⁴ Coherentists can make it rational to believe in P, which fails to cohere with their other beliefs, by destroying those parts of their belief structure that fail to cohere with P. In each case, the *intention* to produce a false sense of justification is, of course, aiming for an epistemically irrational outcome, but it could produce a state of rational belief. Still, this is just an implication of the familiar fact that one can be in a doxastic situation where it is rational to believe what is false.

What if I face a subtler problem: it is irrational and irresponsible to think that my unconfirmed peers pay less attention to the issues than I do, but I *fail to notice* that it is irrational and irresponsible? (If I am an evidentialist, for instance, perhaps I fail to notice some of my evidence that rules this hypothesis out, or, if I am a coherentist, I fail to realize that the hypothesis does not cohere with my other commitments.) In that case, the Earn-a-Spine View does not, of course, endorse my failure to notice that my view is irrational and irresponsible – it only gives me permission to think I am in an epistemically superior position if am *epistemically justified* in thinking my unconfirmed peers pay less attention than I do. But given that I have already made an irrational mistake, it will *seem to me* that the Earn-a-Spine View gives me permission to think I am epistemically superior to my unconfirmed peers.

This distasteful result is ameliorated by two considerations. First, the Earn-a-Spine View is, in this respect, still better than the Presumption in Favor of Self-Trust. Whereas the Presumption in Favor of Self-Trust lets me think myself superior to my unconfirmed peers *on the basis of uncertainty alone*, the Earn-a-Spine View requires that I at least identify my reason for thinking myself superior. There is a fighting chance that I will discover that my unconfirmed peers (or at least many of them) *do* pay as much attention to the issues as I do, and so force me to reconsider my judgment that I am their epistemic superior if I adopt the Earn-a-Spine View. The Presumption of Self-Trust will not demand that I give equal weight to their views as long as there is substantial room for doubt about how credible they are overall.

Second, any internalist view that gives any kind of directions or recommendations – including the EWV itself – faces situations in which someone will be put in a situation to follow those directions or recommendations precisely *by some irrational error*. Following such directions or recommendations cannot, then, be expected to give us an ideally rational outcome (cf. Christensen 2011, 4–5). It is always possible, for instance, that I irrationally think the person I disagree with is my epistemic peer, when, in fact, she is my superior. In such a case, the EWV tells me to split the difference between her view and my own. This result will be sub-optimal because of my previous error in assessing peerhood, but, given what is present to my attention, the adjustment will seem to be rational.

Finally, the Gibbard–Elgin–Foley-style argument in favor of trusting judgments has an odd result: it tells me that I am *more* justified in thinking P is

true if I am *ignorant* of other people's belief that $\sim P$, and that, if I am not very credible, but I do not have evidence of my lack of credibility, I am more justified in trusting my judgments than I would be if I learned about my own lack of credibility. According to my view, I am at least somewhat justified in thinking I am good at grammar if I mistakenly believe I am doing very well on a grammar test, and am unaware of widespread overconfidence biases, and then would become less justified after gaining new information about these biases. On the one hand, it seems paradoxical that those who are more ignorant might be more justified in their confidence. But, on the other hand, it is a familiar observation that intellectual humility is often a result of learning; it makes sense that a self-correcting system will become more aware of its capacity for error over time.¹⁵ And, if the Gibbard–Elgin–Foley argument is correct, this sort of trust in thus-far-unchallenged judgments is the only alternative to bleak skepticism. If we refrain from trusting our judgments until we can show them reliable, and we must rely on our judgments to determine whether our judgments are reliable, we cannot trust in anything at all.

Since there is always a possibility of error, there is always a chance that further epistemic adjustments I make, however rational they appear to me, may be thrown off because of my previous errors. This is a sad fact of human fallibility. But the Earn-a-Spine View does not fare worse, in this regard, than other views that give recommendations to epistemic agents appealing to their first-person point of view, rather than externalist elements that they cannot identify from their own point of view.

IX. Conclusion

Many of our cherished moral, religious, political, and philosophical views are controversial, and many smart people disagree with us, so we should not think our views are *obviously* right. But we should not be too quick to suppose that those who disagree with us are our epistemic peers. Nor should we be too quick to dismiss the possibility. The EWV does not demand that we spinelessly suspend judgment about our cherished controversial beliefs, since people who disagree about these topics are apt to have different credentials – that is, different evidence and different intellectual abilities – which make it hard to determine who is more credible. But if it is an open question whether we are more credible than those who disagree with us, our doxastic attitudes should reflect that uncertainty somehow. Mushy credences might be a way to reflect that uncertainty. But the Earn-a-Spine View provides a less contentious suggestion; we can justify remaining confident in our positions just so long as we can come up with reasonable grounds for thinking ourselves epistemically advantaged according to epistemic standards we don't have reason to think our opponents reject. Personal biases being what they are, this policy will predictably lead to people frequently underestimating rival positions, but it makes them work for

it. Once our grounds are made explicit, there is a much better chance we can either advance the debate by contributing new arguments and evidence, or discover what mistakes we have made and revise our views.

Notes

1. While King cites passages from Kelly (2005, 174–175), Christensen (2009, 756–757) and Elgin (2010, 53), even these passages do not quite match up with his conditions.
2. As King defines peerhood, peers must have the same evidence, but need only have equally good dispositions.
3. King splits the difference, allowing that peers can be equal overall in terms of dispositions, but requiring equality (or sameness) in each of his two conditions.
4. What of comparably influential representatives of the EWW? Feldman's (2007) definition of peerhood lists credibility-conferring features, but leaves it ambiguous whether people must be equal in *each* respect, or equal *overall*: 'people are *epistemic peers* when they are roughly equal with respect to intelligence, reasoning powers, background information, etc.' (Feldman 2007, 201). Christensen (2007, 188–189) describes a clear-cut example of epistemic peers, in which they are the same in *each* of many respects, but he does not actually *define* 'epistemic peer'.
5. Christensen (2011, 16) seems to take a similar approach.
6. Cf. White (2010). See Joyce (2010) for a response.
7. This argument is adapted from Elga (2007, 486–488). The term 'bootstrapping' comes from Vogel (2000, 615).
8. The most straightforward evidence for this claim comes from research on the 'bias blind spot' (Pronin, Lin, and Ross 2002; Pronin 2007), which shows that we tend to overestimate the biases of those who disagree with us, and underestimate our own biases. Even apart from our tendency to impute biases to others, I expect we tend to underestimate how much is going on in other people's minds, and take them to be less responsive to evidence than they actually are, given the prevalence of the Fundamental Attribution Error (cf. Ross and Nisbett 1991) (or something like it), and our tendency to consider only our own strategies for achieving success, while neglecting others' strategies (cf. Kahnemann 2011, 259–264). Further research suggests we have a widespread tendency to think members of a disagreeing *group* – e.g. members of a competing political party or rival religious group – are less rational than our own group (Kenworthy and Miller 2002; Kenworthy 2003; O'Brien and McGarty 2009; Bäck et al. 2010). While I would not want to assume that any of these psychological tendencies is active in every instance of disagreement, or leads us to think we are epistemically superior in every disagreement, they do suggest that, in most disagreements, we at least *think* we have an epistemic advantage.
9. There is one possible exception: if you have discovered that your impressions of epistemic superiority have a good reliable track record, then your impression of superiority, plus your evidence that this impression is reliable, would justify you in concluding that you enjoy an epistemic advantage.
10. Again, info about your own track record could get you around this necessary condition. See fn. 10.
11. If we do not take an all-or-nothing view of belief, our attitude to the explanation that earns us a spine might not be full-fledged belief; we might adopt some

lesser credal state, like thinking the explanation is more likely true than not, or supposing it is true, where 'supposing' involves some level of credence below full belief. Any such credal state would still need to be justified according to standards of rationality and responsibility.

12. Feldman (2006, 223–224) argues that it 'is tenacious and stubborn, but not reasonable' for me to trust *my* insight rather than yours, but he draws this conclusion from the supposition that I know your insight constitutes 'comparable evidence [to my own] supporting a competing position.' Conee (2009, 318–319) discusses cases in which we become confident because we have reasons not shared with our peers, but we *should* expect them to have comparably good reasons they have not shared with us. Conee's conclusion, however, depends on the disagreement in question being a longstanding disagreement among experts in a field. I have some doubts whether Conee's conclusion is decisive even within that context, but I grant that in his scenario there are many factors weighing against having *much* confidence that an unshared insight would tip the balance far toward our side.
13. Elgin and Foley disagree, of course, about the extent to which your judgments are trustworthy in isolation. See section II above.
14. As a reviewer points out, this result can be avoided by a particular form of evidentialists, specifically those holding that my belief that P is justified if it is supported by the evidence I *ought to have*, for some (presumably epistemic) sense of *ought*. Since getting rid of some of my evidence seems likely to be contrary to my epistemic duty, such an evidentialist would *not* need to admit that beliefs resulting from such a breach of duty are justified. I am somewhat doubtful, though, that such an evidentialist theory would still be fully internalist. Suppose I viciously destroy some of my current knowledge – perhaps knowledge about my involvement in a robbery. After the knowledge is destroyed, this form of epistemology would judge my subsequent self to be unjustified in believing I had never committed a robbery, because this belief will be contradicted by evidence my future self *should have had*. But this judgment clearly involves evidence that is not, in any sense, internal to my future self's mind, consciousness, or available intellectual resources.
15. As Ballantyne (2015) points out, the literature on the 'bias blind spot' (Pronin, Lin, and Ross 2002; Pronin 2007) calls into question our ability to tell when those who disagree with us are more biased than ourselves (at least most of the time). But what about those who do not know about the bias blind spot? It is unclear whether they are being irrational in thinking others are more biased – if so, then perhaps they are not justified in thinking others less credible anyway. But it seems clearly true that those who know the empirical findings on the bias blind spot will be *more* informed about their own unreliability, and so will have less of an excuse than those who are ignorant of these findings.

Notes on contributor

Ben Sherman currently holds a position as a lecturer in Philosophy at Brandeis University. In recent years, he has also taught at Boston University, Northeastern University, and Emerson College. His primary research interests are in ethics, epistemology, and the intersection between the two fields.

References

- Bäck, Emma, Peter Esaiasson, Mikael Gilljam, and Torun Lindholm. 2010. "Biased Attributions Regarding the Origins of Preferences in a Group Decision Situation." *European Journal of Social Psychology* 40: 270–281. doi:10.1002/ejsp.618.
- Ballantyne, Nathan. 2015. "Debunking Biased Thinkers (Including Ourselves)." *Journal of the American Philosophical Association* 1 (1): 141–162. doi:10.1017/apa.2014.17.
- Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116 (2): 187–217. doi:10.1215/00318108-2006-035.
- Christensen, David. 2009. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4 (5): 756–767. doi:10.1111/j.1747-9991.2009.00237.x.
- Christensen, David. 2011. "Disagreement, Question-begging and Epistemic Self-criticism." *Philosophers' Imprint* 11 (6): 1–22. <http://hdl.handle.net/2027/spo.3521354.0011.006>.
- Conee, Earl. 2009. "Peerage." *Episteme* 6 (03): 313–323. doi:10.3366/E1742360090000732.
- Elga, Adam. 2007. "Reflection and Disagreement." *NOÛS* 41 (3): 478–502. doi:10.1111/j.1468-0068.2007.00656.x.
- Elgin, Catherine Z. 1996. *Considered Judgment*. Princeton, NJ: Princeton University Press.
- Elgin, Catherine Z. 2010. "Persistent Disagreement." In *Disagreement*, edited by Richard Feldman and Ted Warfield, 53–68. New York: Oxford University Press.
- Feldman, Richard. 2006. "Epistemological Puzzles about Disagreement." In *Epistemology Futures*, edited by Stephen Hetherington, 216–237. New York: Clarendon Press.
- Feldman, Richard. 2007. "Reasonable Religious Disagreements." In *Philosophers without Gods: Meditation on Atheism and the Secular Life*, edited by Louise Antony, 194–214. New York: Oxford University Press.
- Foley, Richard. 2001. *Intellectual Trust in Oneself and Others*. New York: Cambridge University Press.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Gutting, Gary. 1982. *Religious Belief and Religious Skepticism*. Notre Dame, IN: University of Notre Dame Press.
- Joyce, James M. 2010. "A Defense of Imprecise Credences in Inference and Decision Making." *Philosophical Perspectives* 24: 281–323. doi:10.1111/j.1520-8583.2010.00194.x.
- Kahnemann, Daniel. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Kelly, Thomas. 2005. "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology*, Vol. 1, edited by Tamar Szabo Gendler and John Hawthorne, 167–196. Oxford: Clarendon Press.
- Kenworthy, Jared B. 2003. "Explaining the Belief in God for Self, in-group, and out-group Targets." *Journal for the Scientific Study of Religion* 42 (1): 137–146. doi:10.1111/1468-5906.00167.
- Kenworthy, Jared B., and Norman Miller. 2002. "Attributional Biases about the Origins of Attitudes: Externality, Emotionality and Rationality." *Journal of Personality and Social Psychology* 82 (5): 693–707. doi:10.1037/0022-3514.82.5.693.
- King, Nathan. 2012. "Disagreement: What's the Problem? Or a Good Peer is Hard to Find." *Philosophy and Phenomenological Research* 85 (2): 249–272. doi:10.1111/j.1933-1592.2010.00441.x.
- Lackey, Jennifer. 2010. "A Justificationist View of Disagreement's Epistemic Significance." In *Social Epistemology*, edited by Adrian Haddock, Alan Millar, and Duncan Pritchard, 298–325. New York: Oxford University Press.
- O'Brien, Léan V., and Craig McGarty. 2009. "Political Disagreement in Intergroup Terms: Contextual Variation and the Influence of Power." *British Journal of Social Psychology* 48: 77–98. doi:10.1348/014466608X299717.

- Pronin, Emily. 2007. "Perception and Misperception of Bias in Human Judgment." *Trends in Cognitive Sciences* 11 (1): 37–43. doi:[10.1016/j.tics.2006.11.001](https://doi.org/10.1016/j.tics.2006.11.001).
- Pronin, Emily, Daniel Y. Lin, and Lee Ross. 2002. "The Bias Blind Spot: Perceptions of Bias in Self versus Others." *Personality and Social Psychology Bulletin* 28: 369–381. doi:[10.1177/0146167202286008](https://doi.org/10.1177/0146167202286008).
- Ross, Lee, and Richard E. Nisbett. 1991. *The Person and the Situation*. New York: McGraw-Hill.
- Vogel, Jonathan. 2000. "Reliabilism Leveled." *The Journal of Philosophy* 97 (11): 602–623. doi:[10.2307/2678454](https://doi.org/10.2307/2678454).
- White, Roger. 2010. "Evidential Symmetry and Mushy Credence." In *Oxford Studies in Epistemology* Vol. 3, edited by Tamar Szabo Gendler and John Hawthorne, 274–293. New York: Oxford University Press.